

UNITED STATES PATENT APPLICATION FOR:

**SYSTEM AND METHOD FOR RESERVING AND MANAGING MEMORY
SPACES IN A MEMORY RESOURCE**

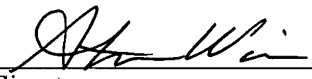
INVENTORS:

**HENRY P. MORETON
JOHN ERIK LINDHOLM
MATTHEW N. PAPAKIPOS
RUI M. BASTOS**

ATTORNEY DOCKET NUMBER: NVDA/P000572

CERTIFICATION OF MAILING UNDER 37 C.F.R. 1.10

I hereby certify that this New Application and the documents referred to as enclosed therein are being deposited with the United States Postal Service on December 2, 2003, in an envelope marked as "Express Mail United States Postal Service", Mailing Label No. EL980273345US, addressed to: Commissioner for Patents, Box PATENT APPLICATION, Alexandria, VA 22313-1450.



Signature

Stephanie Winner

Name

December 2, 2003

Date of signature

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] The application is a continuation-in-part of, and claims priority benefit of, co-pending U.S. patent application Serial No. 10/419,524, titled "System and Method for Reserving and Managing Memory Spaces in a Memory Resource", filed April 21, 2003. The subject matter of the related patent application is hereby incorporated by reference.

FIELD OF THE INVENTION

[0002] One or more aspects of the invention generally relate to computer graphics, and more particularly to reserving, accessing, and managing memory spaces for use by threads executing on a graphics processor.

DESCRIPTION OF THE BACKGROUND

[0003] Conventionally, graphics data is processed on a graphics processor through the use of threads executing on the graphics processor. Threads executing on the graphics processor generate thread data, such as, source data, destination data, and intermediate data generated during execution of the thread. The thread data is typically stored into a first available memory location of a memory resource being used to store thread data. Therefore, thread data generated by a same thread may be located in non-neighboring locations throughout the memory resource.

[0004] Conventionally, memory resources storing thread data in non-neighboring locations produce significant cross-thread interaction where one thread inadvertently accesses thread data of another thread. Memory resources storing thread data in non-neighboring locations also produce low access coherency where thread data in neighboring memory locations are unrelated. Access coherency is important in graphics processing since a single read command for a memory location typically accesses the memory location and a group of neighboring memory locations in one retrieval. Therefore, when a read command is performed, it is desirable to have high access coherency so that the group of neighboring memory locations that is read contain related thread data.

SUMMARY

- [0005]** New systems and methods for reserving a memory space in a memory resource for each thread in a first set of threads are disclosed.
- [0006]** A method for memory allocation for a multithreaded processor includes obtaining threads. Sets are created from the threads according to thread type. Memory space for each thread in a set is allocated responsive to the thread type of the set.
- [0007]** A method for memory allocation for a multithreaded processor includes obtaining threads. Sets are created from the threads according to thread type. Memory space for each thread in a set is allocated responsive to the thread type of the set. Memory space for data buffers which is accessible by the threads is allocated.
- [0008]** A method for accessing a memory resource for a plurality of threads simultaneously executable in a graphics processor responsive to a graphics program module includes receiving a first sample to be processed by a first thread and receiving a second sample to be processed by a second thread. A first memory space is accessed only by the first thread during processing of the first sample by the first thread. A second memory is accessed only by the second thread during processing of the second sample by the second thread. A third memory space is accessed by the first thread and the second thread.
- [0009]** A computer program product having a computer readable medium having computer program instructions recorded thereon, said computer program product includes instructions for determining a first set of threads of a plurality of threads, each thread of said plurality of threads being associated with a graphics program module executing on a graphics processor, instructions for allocating a first memory space in a memory resource to each thread in the first set of threads, the first memory space being reserved for the thread to which the first memory space is allocated, and instructions for allocating a second memory space in the memory resource to a data buffer, the data buffer being accessible by each thread in the plurality of threads.
- [0010]** Another computing system includes a memory resource, a graphics processor coupled to the memory resource for executing one or more graphics program modules

and a central processing unit (CPU) coupled to the memory resource and the graphics processor. The CPU or the graphics processor determines a first set of threads from a plurality of threads simultaneously executable in said graphics processor, allocates a memory space to be each thread in said first set of threads to respectively reserve the memory space. The CPU allocates an additional memory space accessible by each thread in the first set of threads to one or more data buffers to respectively reserve the additional memory space.

BRIEF DESCRIPTION OF THE VARIOUS VIEWS OF THE DRAWINGS

[0011] Accompanying drawing(s) show exemplary embodiment(s) in accordance with one or more aspects of the invention; however, the accompanying drawing(s) should not be taken to limit the invention to the embodiment(s) shown, but are for explanation and understanding only.

[0012] FIG. 1 is a block diagram illustrating a computing system;

[0013] FIG. 2 is a block diagram of a memory resource containing memory spaces that are reserved, accessed, or managed according to a memory space reservation process;

[0014] FIG. 3A shows a flowchart for reserving memory spaces for threads of a first set of threads executing on a graphics processor;

[0015] FIG. 3B shows a flowchart for accessing memory spaces that have been reserved for threads of a first set of threads using the flow shown in FIG. 3A;

[0016] FIGS. 4A and 4B show flowcharts for managing memory spaces for use by threads of a first set of threads executing on a graphics processor;

[0017] FIG. 5 is a block diagram of a memory resource containing memory spaces that are reserved, accessed, or managed according to another memory space reservation process;

[0018] FIG. 6 is a block diagram of a memory resource containing memory spaces that are reserved, accessed, or managed according to a first alternative memory space reservation process to that of FIG. 5;

[0019] FIG. 7 is a block diagram of a memory resource containing memory spaces that are reserved, accessed, or managed according to a second alternative memory space reservation process to that of FIG. 5;

[0020] FIG. 8 is a block diagram of a memory resource for storing thread data from a primary memory resource;

[0021] FIG. 9 shows a flowchart for transferring data between a primary memory space of a primary memory resource (shown in FIG. 8) and a memory space of a memory resource;

[0022] FIG. 10A is a conceptual diagram of a data buffer used by a display device.

[0023] FIG. 10B is a memory resource containing memory spaces that are reserved, accessed, or managed according to a memory space reservation process;

[0024] FIG. 11A shows a flowchart for reserving memory spaces for one or more data buffers and threads of a first set of threads executing on a graphics processor;

[0025] FIG. 11B shows a flowchart for accessing memory spaces that have been reserved for one or more data buffers and threads of a first set of threads using the flow shown in FIG. 11A; and

[0026] FIG. 11C shows another flowchart for accessing memory spaces that have been reserved for one or more data buffers and threads of a first set of threads using the flow shown in FIG. 11A.

DETAILED DESCRIPTION

[0027] FIG. 1 is a block diagram of a Computing System 100 in which embodiments in accordance with one or more aspects of the invention may be used. Computing System 100 may be a desktop computer, server, laptop computer, palm-sized computer, tablet computer, game console, cellular telephone, computer-based simulator, or the like. Computing System 100 includes a Host Computer 110, a Graphics Subsystem 120, an "External" Bus 116, and a Display 185. Graphics Subsystem 120 contains a Graphics Processor 125 and a Local Memory 135. By "External Bus," it is meant a bus used to put Host Computer 110 in communication with a subsystem not part of Host Computer 110.

[0028] Graphics Processor 125 contains Functional Units 140, 150, 160, and 170, a Graphics Interface 117, a Thread Control Buffer (TCB) 127, an Address Unit (AU) 128, a Memory Controller 130, and a Scanout 180. In an alternative embodiment, Functional Units 140, 150, 160, and 170, each contain a Local Storage Resource (LSR) 145, 155, 165, and 175, respectively. Though a pipeline architecture is shown for Graphics Processor 125, Graphics Processor 125 may be implemented as a multi-processor architecture where Functional Units 140, 150, 160, 170 are not cascaded.

[0029] TCB 127 may be located in Memory Controller 130 (as shown), in a Functional Unit 140, 150, 160, or 170, in Graphics Interface 117, or located as a separate device coupled to Memory Controller 130 and a Functional Unit 140, 150, 160, or 170. Likewise, AU 128 may be located in Memory Controller 130 (as shown), in a Functional Unit 140, 150, 160, or 170, in Graphics Interface 117, or located as a separate device coupled to Memory Controller 130 and a Functional Unit 140, 150, 160, or 170. In an alternative embodiment, TCB 127 and AU 128 are integrated on the same chip. One or more embodiments, in accordance with one or more aspects of the invention, to configure TCB 127 and AU 128 include computer program products having a computer readable medium with computer program instructions to perform particular functions of such embodiments.

[0030] Host Computer 110 communicates with Graphics Subsystem 120 via External Bus 116 and Graphics Interface 117. Host computer 110 includes a Host Memory 112, a Host Central Processing Unit (CPU) 114, and a System Interface 115. Host CPU 114 may include a system memory controller to interface directly to Host Memory 112. Or Host CPU 114 may interface with System Interface 115 and communicate with Host Memory 112 through System Interface 115. System Interface 115 may be an I/O (input/output) interface or a bridge device including the system memory controller to interface directly to Host Memory 112. Examples of System Interface 115 known in the art include Intel ® Northbridge and Intel ® Southbridge.

[0031] Host Computer 110 loads a Graphics Application 111 into Host Memory 112 and Host CPU 114 executes program instructions of Graphics Application 111. Graphics Application 111 generates streams of graphics data used to generate an image to be displayed on Display 185. Alternatively, the streams of data are output via Graphics Interface 117 to a film recording device or written to a peripheral device, e.g., disk drive, tape, compact disk, or the like. Program instructions of Graphics Application 111, subsets of program instructions of Graphics Application 111 (i.e., programs or graphics program modules), and graphics data are read from or stored to a memory resource, e.g., any combination of Host Memory 112, Local Memory 135, and LSR 145, 155, 165, or 175. When a portion of Host Memory 112 is used to store program instructions and graphics data, the portion of Host Memory 112 can be uncached so as to increase performance of access by Graphics Processor 125.

[0032] When Graphics Application 111 is loaded to and executed by Host CPU 114, an operating system executing in Host CPU 114 calls a Driver 113. Driver 113 then executes in Host CPU 114. As used herein, Driver 113 is a program interface between Graphics Application 111 and Graphics Processor 125. Driver 113 executes when Graphics Application 111 is active. Driver 113 may include a computer program product having a computer readable medium that includes computer program instructions configured to perform particular functions of one or more embodiments in accordance with one or more aspects of the invention.

[0033] Host Computer 110 communicates the graphics data used to generate an image to Graphics Processor 125 via System Interface 115, External Bus 116, and Graphics Interface 117. The graphics data generated by Graphics Application 111 may include a high-level description of a scene for display, and other high-level information such as from where the scene is to be viewed, what textures should be applied to different primitives in the scene, and where lights are located in the scene. Display 185, however, may be a relatively simple device for accepting and outputting color information on a pixel-by-pixel basis that cannot interpret the high-level graphics data from Graphics Application 111.

[0034] Therefore, the high-level graphics data is processed by Functional Units 140, 150, 160, and 170 of Graphics Processor 125 and translated into pixel color information for the image to be displayed on Display 185. Functional Units 140, 150, 160, and 170 of Graphics Processor 125 are programmable units capable of executing programs of Graphics Application 111. In one embodiment, a Functional Unit 140, 150, 160, or 170 is a programmable vertex processor capable of performing per-vertex computations (such as lighting and time-varying spatial offsets), subdivision surface algorithms (as known in the art), and N-patch algorithms (or “normal patch”, as known in the art). In another embodiment, a Functional Unit 140, 150, 160, or 170 is a programmable shader processor capable of performing per-fragment operations, such as texturing, lighting, bump mapping, or the like.

[0035] The graphics data processed by Functional Units 140, 150, 160, and 170 of Graphics Processor 125 can be primitive data, surface data, pixel data, vertex data, fragment data, or the like. For simplicity, the remainder of this description will use the term “sample” to refer to primitive data, surface data, pixel data, vertex data, fragment data, or the like. The number of Functional Units 140, 150, 160, and 170 are for illustrative purposes only and Graphics Processor 125 may include more or fewer Functional Units without departing from the spirit or scope of one or more aspects of the invention.

[0036] A sample and sample type identifier associated with the sample are received by a Functional Unit 140, 150, 160, or 170. The sample type identifier associated with a

sample identifies the sample type of the sample (e.g., pixel sample, vertex sample, fragment sample, or the like.) and determines which particular Functional Unit 140, 150, 160, or 170 receives and processes the sample. For example, if a sample type identifier identifies a sample as a vertex sample and Functional Unit 140 is a vertex processor, the vertex sample is received and processed by Functional Unit 140.

[0037] Along with a sample and a sample type identifier associated with the sample, a Functional Unit 140, 150, 160, or 170 receives a pointer to a program associated with the sample. A program (or graphics program module) is a specific subset of program instructions of Graphics Application 111 used to process the associated sample. The program associated with a sample relates to the sample type associated with the sample. For example, if a sample is a vertex sample, the vertex sample is processed according to a vertex program, i.e., a program configured to process vertex samples. The pointer to the associated program locates a memory address in a memory resource (e.g., Host Memory 112, Local Memory 135, LSR 145, 155, 165, or 175, or the like.) where the associated program is found. A same program having a same associated pointer can be used to process several different samples. In an alternative embodiment Memory Controller 130 contains a Cache 129 for caching graphics data and program instructions read from Local Memory 135 or Host Memory 112.

[0038] When Functional Unit 140, 150, 160, or 170 receives a sample, a sample type identifier associated with the sample, and a pointer to a program associated with the sample, a thread is assigned to the sample by TCB 127. Graphics Processor 125 can execute a predefined number of threads in parallel. TCB 127 includes storage resources to retain thread state data to track the number of threads previously assigned to other samples and the number of threads still available. In an alternative embodiment, the storage resources of TCB 127 also contain a thread type identifier for each thread to enable TCB 127 to assign a thread to a sample based on the thread type identifier of the thread and the sample type identifier of the sample.

[0039] As used herein, a thread is a set of processes for processing a sample according to a program associated with the sample. A thread assigned to a sample uses a pointer to locate the program associated with the sample and loads and executes the

program on the Functional Unit 140, 150, 160, or 170 that received the sample. The thread then processes the sample according to the associated program executing on the Functional Unit 140, 150, 160, or 170 that received the sample. Graphics Processor 125 contains one or more Functional Units 140, 150, 160, or 170 and can process one or more threads simultaneously, each thread executing an associated program (or graphics program module) on a Functional Unit 140, 150, 160, or 170.

[0040] Each thread of a plurality of threads executing on Graphics Processor 125 uses a predefined amount of memory space for storing thread data generated by the thread during its execution. Thread data to be stored in memory space include, for example, source data, destination data, and intermediate data generated during execution of the thread. A memory space used for storing thread data can be located in Host Memory 112, a peripheral memory resource (not shown) coupled to System Interface 115 (e.g., hard drive, Zip drive, tape drive, CD-R, CD-RW, etc.), a graphics memory resource, such as Local Memory 135 or LSR 145, 155, 165, or 175, or any combination of the above. Memory spaces in a memory resource for use by threads of Graphics Processor 125 are allocated and managed in accordance with one or more aspects of the invention as described in relation to exemplary embodiments illustratively shown in FIGS. 3A, 3B, 4A, and 4B.

[0041] In an alternative embodiment, threads executing on Graphics Processor 125 can be of different thread types configured to process samples of different sample types. Threads of different thread types may use different memory space amounts for storing thread data. A thread type of a thread is identified by a thread type identifier and a sample type of a sample is identified by a sample type identifier. TCB 127 assigns threads of a particular thread type to match samples of a particular sample type by matching the thread type identifier of the thread with the sample type identifier of the sample. For example, a vertex sample is processed by a vertex thread and a fragment sample is processed by a fragment thread, the vertex thread may use a greater memory space amount than the fragment thread. Furthermore, another sample type, such as a primitive sample is processed by a primitive thread, the primitive thread may use a different memory space amount than either the fragment thread or the vertex thread.

[0042] In an alternate embodiment, threads of a same thread type may use different memory space amounts for storing thread data. A memory space amount for a thread is dynamically allocated by TCB 127 prior to execution of the thread. A thread memory space is specified for each thread by Graphics Application 111 and is included in the program instructions output to Graphics Processor 125.

[0043] As stated above, a thread processes a sample according to a program associated with the sample, the program executing on the Functional Unit 140, 150, 160, or 170 that received the sample. During the processing of the sample, the thread executes an instruction in the program that generates an access command (i.e., read, or write command) for a memory space reserved for the thread. If the memory space reserved for the thread is implemented as a stack, the access command can be referred to as a “push” or “pop” command. Memory spaces in a memory resource for use by threads of Graphics Processor 125 are reserved or managed in accordance with one or more aspects of the invention as described in relation to exemplary embodiments illustratively shown in FIGS. 3A, 3B, 4A, and 4B.

[0044] An access command produced by a thread executing on a Functional Unit 140, 150, 160, or 170 is sent to AU 128 for further processing. An access command includes an operation command (i.e., read or write) and address request information. Address request information is used by AU 128 in determining the memory location address of a memory location in memory space that the access command is to access in accordance with one or more aspects of the invention as described in relation to exemplary embodiments illustratively shown in FIGS. 3A, 3B, 4A, and 4B.

[0045] AU 128 includes an AU computational unit and an AU storage unit. The AU computational unit determines memory location addresses for received access commands using a look-up table or a predefined computation. The AU computational unit can be hard-wired to perform these functions or be configured by a software program to perform the functions. The AU storage unit is used to store memory location address computation information needed by the AU computational unit to determine memory location addresses for received access commands. The AU storage unit and the AU computational unit may be on separate chips or integrated on the same chip. Memory

location addresses of memory locations in a memory space allocated to a thread are determined in accordance with one or more aspects of the invention as described in relation to exemplary embodiments illustratively shown in FIGS. 3A, 3B, 4A, and 4B.

[0046] After AU 128 determines a memory location address for a received access command, the memory location address and operation command are sent to Memory Controller 130. Memory Controller 130 arbitrates between hardware components of Graphics Subsystem 120 initiating access commands to memory resources containing memory spaces used by threads executing on Graphics Processor 125. Examples of such memory resources are Host Memory 112, a peripheral memory resource (not shown) coupled to System Interface 115 (e.g., hard drive, Zip drive, tape drive, CD-R, CD-RW, or the like), a graphics memory resource, such as Local Memory 135 or LSR 145, 155, 165, or 175, or any combination of the above. Memory Controller 130 receives the memory location address and operation command and accesses the memory location identified by the memory location address according to the operation command.

[0047] In an alternate embodiment AU 128 identifies and avoids read-after-write (RAW) hazards using a method known in the art. RAW hazards can occur when write operations are coalesced such that order is not maintained between read and write operation commands for each memory location received by Memory Controller 130 from AU 128. For example, a RAW hazard occurs when a coalesced write to a memory location is delayed such that a read from the memory location occurs before instead of after the write to the memory location. Likewise, RAW hazards can occur when order is not maintained between read and write operation commands for each memory location received by Local Memory 135 or Host Memory 112 from Memory Controller 130.

[0048] FIG. 2 shows a block diagram of a Memory Resource 200 containing Memory Spaces 205 that are reserved, accessed, or managed according to a memory space reservation process. FIG. 2 is described in relation to FIG. 1.

[0049] Memory Resource 200 can be Host Memory 112, a peripheral memory resource (not shown) coupled to System Interface 115 (e.g., hard drive, Zip drive, CD tape drive, CD-R, CD-RW, or the like), a graphics memory resource (such as Local Memory 135 or LSR 145, 155, 165, or 175), or any combination of the above. Memory

Resource 200 or portions of Memory Resource 200 (e.g., Memory Space 205) can be implemented in various configurations, for example, as a stack, cache, FIFO, or to support random access. Memory Resource 200 or portions of Memory Resource 200 (e.g., Memory Space 205) can be of different memory types, for example, a set of registers, RAM, DRAM, or the like. Memory Resource 200 can be internal to Graphics Processor 125 (i.e., located on the same chip as Graphics Processor 125) such as LSR 145, 155, 165, or 175. Memory Resource 200 can also be external to Graphics Processor 125 (i.e., not located on the same chip as Graphics Processor 125), such as Host Memory 112, Local Memory 135, or a peripheral memory resource. Additionally, Memory Resource 200 can include a memory resource internal to Graphics Processor 125 and a memory resource external to Graphics Processor 125.

[0050] Memory Resource 200 contains at least two Memory Spaces 205 for use by threads of a first set of threads executing on Graphics Processor 125. The first set of threads can include all threads executing on Graphics Processor 125 or a subset of all threads executing on Graphics Processor 125. A Memory Space 205 is reserved for each thread in the first set of threads. Each thread in the first set of threads is identified by a Thread Identification Number (THD#) 202, the THD# 202 of a thread in the first set of threads being an order number of the thread in the first set of threads. In the example shown in FIG. 2, THD#s 202 range from 0 through N1 and the first set of threads includes N1+1 threads.

[0051] Each Memory Space 205 contains at least one Memory Location 210. Each Memory Location 210 is identified by a unique memory location address and has a memory location size or width. A size of a Memory Space 205 (or memory space size) is equal to the number of Memory Locations 210 contained in Memory Space 205 multiplied by the size of a Memory Location 210. The size of a Memory Location 210 can be any number of bits as specified in hardware or software. For simplicity, as used herein, the size of a Memory Location 210 is represented as being equal to 1. Therefore, in the example shown in FIG. 2, the first memory space size of each Memory Space 205 is S1+1.

[0052] A first Memory Location 210 of Memory Space 205 is defined as a Base Memory Space Location 215, the Base Memory Space Location 215 having a unique memory location address. As used herein, the memory location address of Base Memory Space Location 215 is also referred to as a base memory space address of Memory Space 205. If Memory Space 205 is implemented as a stack, thread data is accessed from Memory Space 205 starting from Base Memory Space Location 215 of Memory Space 205. In other words, thread data is stored to (“pushed”) and read from (“popped”) a “top” of Memory Space 205.

[0053] Memory Spaces 205 reserved for threads of the first set of threads includes a Memory Section 220. A first Memory Location 210 of Memory Section 220 is defined as a Base Memory Section Location 225, the Base Memory Section Location 225 having a unique memory location address. As used herein, the memory location address of Base Memory Section Location 225 is also referred to as a base memory section address of Memory Section 220. As shown in FIG. 2, Base Memory Section Location 225 of Memory Section 220 is also Base Memory Space Location 215 for the Memory Space 205 reserved for a thread having THD# 202 of 0. A size of Memory Section 220 is equal to the sum of the memory space sizes of Memory Spaces 205 allocated to threads of the first set of threads.

[0054] FIG. 3A shows a flowchart of a method for reserving Memory Spaces 205 for threads of a first set of threads executing on Graphics Processor 125. FIG. 3A is described in relation to FIGS. 1 and 2.

[0055] At step 300, Host Computer 110 loads Graphics Application 111 into Host Memory 112 and Host CPU 114 executes program instructions of Graphics Application 111. Program instructions of Graphics Application 111, programs or graphics program modules of Graphics Application 111 to be executed by Graphics Processor 125, and graphics data generated by Graphics Application 111 are read from or stored to a memory resource, e.g., any combination of Host Memory 112, Local Memory 135, and LSR 145, 155, 165, or 175. In an alternate embodiment program instructions of Graphics Application 111 are executed directly by Graphics Processor 125.

[0056] At step 305, an operating system executing in Host CPU 114 calls Driver 113 and Driver 113 executes in Host CPU 114. As stated above, Driver 113 is a program interface between Graphics Application 111 and Graphics Processor 125. Driver 113 determines information relating to Graphics Processor 125 including a number of threads Graphics Processor 125 is capable of executing simultaneously, an amount of memory space used by each thread, and a size of memory space available in graphics memory resources (such as Local Memory 135 and LSR 145, 155, 165, or 175) located on Graphics Processor 125. The information relating to Graphics Processor 125 may be contained in Driver 113. In a further alternative embodiment, the information relating to Graphics Processor 125 is hard-wired in AU 128 or specified in Graphics Application 111 or another program from which Driver 113 or AU 128 can receive the information.

[0057] At step 310, Driver 113 determines a size of Memory Section 220 in Memory Resource 200 to be reserved for a first set of threads that will execute on Graphics Processor 125. To do so, Driver 113 determines a number of threads in the first set of threads which is equal to or less than the number of threads Graphics Processor 125 is capable of executing simultaneously. In the example shown in FIG. 2, the first set of threads includes $N1+1$ threads. In an alternate embodiment, Graphics Processor 125 completes step 310, determining the size of Memory Section 220 in Memory Resource 200 to be reserved for the first set of threads that will execute on Graphics Processor 125.

[0058] Driver 113 also determines a first memory space size of a Memory Space 205 to be reserved for each thread in the first set of threads, the first memory space size being equal to or greater than the amount of memory space used by each thread. In the example shown in FIG. 2, the first memory space size of each Memory Space 205 is $(S1+1)$. The memory space size of a Memory Space 205 is equal to the number of Memory Locations 210 contained in Memory Space 205. Driver 113 then multiplies the number of threads in the first set of threads with the first memory space size to determine the size of Memory Section 220 in Memory Resource 200. In an alternate embodiment, Graphics Processor 125 determines the first memory space size of a Memory Space 205 to be reserved for each thread in the first set of threads.

[0059] At step 312, Driver 113 determines which Memory Resource 200 is to contain Memory Section 220. Examples of Memory Resource 200 candidates are Host Memory 112, a peripheral memory resource (not shown) coupled to System Interface 115 (e.g., hard drive, Zip drive, tape drive, CD-R, CD-RW, or the like), a graphics memory resource, such as Local Memory 135 or LSR 145, 155, 165, or 175, or any combination of the above. If the size of memory space available in graphics memory resources located on Graphics Processor 125 is equal to or greater than the size of Memory Section 220 determined at step 310, Driver 113 may determine that a graphics memory resource (such as Local Memory 135 and LSR 145, 155, 165, or 175) is to contain Memory Section 220. In this case, at step 315, Driver 113 assigns a base memory section address for Memory Section 220. If Driver 113 determines that Host Memory 112 or a peripheral memory resource (not shown) coupled to System Interface 115 is to contain Memory Section 220, Driver 113 receives, at step 315, a base memory section address for Memory Section 220 from Host CPU 114. As used herein, the first memory space size of a Memory Space 205 in Memory Section 220 and the base memory section address for Memory Section 220 are referred to as memory section information for Memory Section 220. In an alternate embodiment steps 312 and 315 are completed by Graphics Processor 125, determining which Memory Resource 200 is to contain Memory Section 220 and assigning the base memory section address for Memory Section 220.

[0060] At step 320, Driver 113 allocates a Memory Space 205 to a thread in the first set of threads, starting with a thread having a THD# 202 of 0. Driver 113 identifies a base memory space address of Memory Space 205 (i.e., a memory location address of a Base Memory Space Location 215 of Memory Space 205) to be allocated to the thread. In an alternative embodiment, at step 320, Driver 113 also identifies a memory location address for each Memory Location 210 in Memory Space 205. The memory location address for a Memory Location 210 in Memory Space 205 is determined from the base memory space address of Memory Space 205 and a memory location offset, such as a memory location offset ranging from 0 to S1. Each Memory Location 210 in Memory Space 205 has an associated memory location offset, Base Memory Space Location 215 having an associated memory location offset of 0. As stated above, as used herein, the size of a Memory Location 210 is represented as being equal to 1. Therefore, in the

example shown in FIG. 2, the memory location offsets range from 0 to S1. In an alternate embodiment, Graphics Processor 125 completes step 320, allocating a Memory Space 205 to the thread in the first set of threads.

[0061] If Memory Section 220 is implemented as a stack, Driver 113 or Graphics Processor 125 identifies the base memory space address for Memory Space 205 allocated to a thread using a first predefined computation. The first predefined computation may be a following predefined equation:

base memory section address of Memory Section 220 + (THD# 202 * first memory space size).

[0062] In an alternative embodiment, Memory Section 220 is implemented to support random access. In this case, Driver 113 or Graphics Processor 125 identifies a memory location address for each Memory Location 210 (including a base memory space address for Base Memory Space Location 215) in Memory Space 205 using a second predefined computation. The second predefined computation may be a following predefined equation:

base memory section address of Memory Section 220 + (THD# 202 * first memory space size) + memory location offset.

[0063] For example, as shown in FIG. 2, a base memory space address of Memory Space 205 reserved for a thread having THD# 202 of 1 is determined by the sum of the base memory section address for Memory Section 220 (the memory location address of Base Memory Section Location 225) and 1 (THD# 202) multiplied by S1+1 (first memory space size). If Memory Section 220 is implemented to support random access, the memory location address of Example Memory Location 240 can be determined by the sum of the base memory space address of Memory Space 205 reserved for the thread

having THD# 202 of 1 (as determined above) and a memory location offset (equal to 1) associated with Example Memory Location 240.

[0064] In an alternative embodiment, the first predefined computation is a following predefined order of concatenation:

{base memory section address for Memory Section 220, THD#202, memory location offset of 0}.

[0065] Prior to concatenation, the base memory section address for Memory Section 220 may be truncated to preserve a number of high bits. The number of bits used to represent each component, e.g., base memory section address for Memory Section 220, memory location offset of 0, THD# 202, in the concatenation is fixed at a number of bits needed to represent the largest possible value for that component. For example the number of bits used to represent memory location offset of 0 and memory location offset is $\log_2(\text{first memory space size})$ and the number of bits used to represent THD#202 is $\log_2(N1+1)$. In a further embodiment, the second predefined computation is a following predefined order of concatenation:

{base memory section address for Memory Section 220, THD#202, memory location offset}.

[0066] At step 325, the THD# 202 of the thread allocated Memory Space 205 (at step 320) and the base memory space address of the allocated Memory Space 205 are added to a look-up table. This step is referred to as adding a Memory Space 205 to the look-up table. The look-up table is used to keep a record of Memory Spaces 205 allocated to threads in the first set of threads.

[0067] At step 325, if the allocated Memory Space 205 is implemented to support random access, a memory location offset associated with each Memory Location 210 in

the allocated Memory Space 205 and a memory location address for each Memory Location 210 are also added to the look-up table. This is also referred to as adding a Memory Space 205 to the look-up table.

[0068] At step 330, Driver 113 determines if there is another thread in the first set of threads to be allocated a Memory Space 205. If so, at step 320, a thread having a next order number (i.e., a next THD# 202) in the first set of threads is allocated a Memory Space 205 by Driver 113. At step 330, the allocated Memory Space 205 is added to the look-up table. In the example shown in FIG. 2, $N1 + 1$ threads in the first set of threads are allocated a Memory Space 205, each Memory Space 205 being reserved for the thread to which it is allocated. In an alternate embodiment, Graphics Processor 125 completes step 330, determining if there is another thread in the first set of threads to be allocated a Memory Space 205.

[0069] If Driver 113 determines that there are no other threads in the first set of threads to be allocated a Memory Space 205, the method proceeds to step 335. At step 335, the look-up table is stored to AU 128.

[0070] In a further embodiment, the information relating to Graphics Processor 125 is stored to AU 128 at step 305. As described above, information relating to Graphics Processor 125 includes a number of threads Graphics Processor 125 is capable of executing simultaneously, an amount of memory space used by each thread, and a size of memory space available in graphics memory resources located on Graphics Processor 125. Steps 310 through 330 are then performed by AU 128 instead of Driver 113. For example, at step 310, AU 128 determines the size of Memory Section 220 reserved for the first set of threads by multiplying the number of threads in the first set of threads with the first memory space size. At step 312, AU 128 determines which Memory Resource 200 is to contain Memory Section 220 and loads the base memory section address for Memory Section 220 from Driver 113 or Host CPU 114. At steps 320 through 330, AU 128 allocates a Memory Space 205 for each thread in the first set of threads, a record of the allocations being stored in a look-up table. AU 128 can be hard-wired to perform the functions used in steps 310 through 330 or be configured by a software program to perform the functions used in steps 310 through 330.

[0071] FIG. 3B shows a flowchart of a method for accessing Memory Spaces 205 that have been reserved for threads of a first set of threads using the method shown in FIG. 3A. FIG. 3B is described in relation to FIGS. 1, 2, and 3A.

[0072] At step 340, a Functional Unit 140, 150, 160, or 170 receives a sample, a sample type identifier associated with the sample, and a pointer to a program associated with the sample. A program is a specific subset of program instructions of Graphics Application 111 (executing on Host CPU 114) used to process an associated sample. The pointer to the associated program locates a memory address in a memory resource (e.g., Host Memory 112, Local Memory 135, LSR 145, 155, 165, 175, or the like) where the associated program is found.

[0073] At step 345, TCB 127 assigns a thread to the sample received by Functional Unit 140, 150, 160, or 170, the thread being identified by a THD# 202. In an alternative embodiment, TCB 127 also assigns a stack pointer to the sample, the stack pointer being associated with the thread assigned to the sample and is used in accessing a memory space implemented as a stack.

[0074] At step 350, the thread assigned to the sample uses the received pointer to locate the program associated with the sample and loads the program to the Functional Unit 140, 150, 160, or 170 that received the sample. In an alternative embodiment, TCB 127 assigns a thread (identified by a THD# 202) and a base memory space address of the allocated Memory Space 205 to the sample.

[0075] At step 355, the thread processes the sample according to the program associated with the sample, the program executing on the Functional Unit 140, 150, 160, or 170 that received the sample. During the processing of the sample, the thread executes an instruction in the program that generates an access command for a Memory Location 210 in a Memory Space 205 allocated to the thread (in steps 320 through 330). If Memory Space 205 allocated to the thread is implemented as a stack, the access command can also be referred to as a “push” or “pop” command. The allocated Memory Space 205 to be accessed is reserved for use by the thread to which it is allocated.

[0076] At step 360, an access command produced by the thread is received at AU 128. In an alternative embodiment, the stack pointer associated with the thread

producing the access command is also received at AU 128. An access command includes an operation command and address request information. Address request information is used by AU 128 in determining a memory location address of a Memory Location 210 in the Memory Space 205 to be accessed. If the operation command is a write command, the access command also includes data to be written. If the Memory Space 205 to be accessed is implemented as a stack, the address request information includes the THD#, 202 that identifies the thread producing the access command. If the Memory Space 205 to be accessed is implemented to support random access, the address request information includes the THD# 202 and a memory location offset provided by the instruction generating the access command. The memory location offset is used to locate a particular Memory Location 210 of the Memory Space 205 to be accessed.

[0077] In an alternative embodiment, clamping is performed at steps 362 and 364. To perform clamping, a maximum memory location offset of a Memory Space 205 is set to equal the memory space size of Memory Space 205 (determined at step 310) minus one. In the example shown in FIG. 2, the first memory space size is equal to $(S1+1)$ and, therefore, the maximum memory location offset for Memory Space 205 is equal to $S1$.

[0078] At step 362, AU 128 determines if a memory location offset received at step 360 is greater than the maximum memory location offset of the Memory Space 205 to be accessed. If so, at step 364, the received memory location offset is set to equal the maximum memory location offset. Otherwise, the method proceeds to step 365.

[0079] The clamping function performed at steps 362 and 364 provides for error correction of illegal access commands, i.e., access commands containing memory location offsets that exceed the maximum memory location offset of a Memory Space 205 to be accessed. In a further embodiment, an error condition indicator is stored in AU 128 and is set if AU 128 determines (at step 362) that an illegal access command has been received. In yet a further embodiment, a received illegal access command (determined at step 362) is ignored if the received illegal access command is a write command and a value of 0 is returned if the received illegal access command is a read command.

[0080] At step 365, AU 128 uses the look-up table stored to AU 128 (at step 335) and the address request information received by AU 128 (at step 360) to determine a memory location address of a Memory Location 210 in the Memory Space 205 to be accessed by the access command (received at step 360). If the Memory Space 205 to be accessed is implemented as a stack, AU 128 uses THD# 202 contained in the address request information to determine the memory location address of Base Memory Space Location 215 (i.e., the base memory space address of Memory Space 205) from the look-up table. At step 365, if the Memory Space 205 to be accessed is implemented to support random access, AU 128 uses THD# 202 and the memory location offset contained in the address request information to determine the memory location address of the Memory Location 210 to be accessed from the look-up table. At step 367, the memory location address of the Memory Location 210 and the operation command is used by AU 128 to determine if a RAW hazard exists. If so, step 367 is repeated until the write operation to the memory location address of the Memory Location 210 is completed. When a RAW hazard does not exist in step 367, the method proceeds to step 370. In a further embodiment, the operation command and address request information received by AU 128 is used to determine if a RAW hazard exists.

[0081] At step 370, the operation command contained in the access command (received at step 360) and the memory location address (determined at step 365) of the Memory Location 210 to be accessed by the access command is sent to Memory Controller 130. If the operation command is a write command, data to be written is also sent to Memory Controller 130. In an alternative embodiment, the stack pointer associated with the thread producing the access command is also sent to Memory Controller 130. Memory Controller 130 then accesses the Memory Location 210 in Memory Space 205 according to the operation command.

[0082] At step 355, the thread continues processing of the sample (received at step 340) according to the program executing on the Functional Unit 140, 150, 160, or 170 that received the sample until the program completes execution. While the thread is processing the received sample, the Memory Space 205 used by the thread is accessed only by that particular thread.

[0083] The method shown in the flowchart of FIG. 3B can be used to access a Memory Space 205 reserved for another thread in the first set of threads as well. The method of FIG. 3B repeats for each thread in the first set of threads that is assigned (at step 345) to a sample by TCB 127.

[0084] FIGS. 4A and 4B show flowcharts of a method for managing Memory Spaces 205 for use by threads of a first set of threads executing on Graphics Processor 125. FIG. 4A is described in relation to FIGS. 1, 2, and 3A. Steps 300 through 315 of FIG. 4A are substantially similar to steps 300 through 315 of FIG. 3A and are not described in further detail. At step 420, the memory section information for Memory Section 220 (i.e., the first memory space size of Memory Space 205 and the base memory section address of Memory Section 220) as determined in steps 310 and 315 is stored to AU 128.

[0085] FIG. 4B is described in relation to FIGS. 1, 2, 3A, and 4A. Steps 340 through 364 of FIG. 4B are substantially similar to steps 340 through 364 of FIG. 3B and are not described in further detail.

[0086] At step 465, memory location address computation information is loaded into the AU computational unit of AU 128. As used herein, memory location address computation information refers to information used by AU 128 to determine a memory location address of a Memory Location 210 in a Memory Space 205 to be accessed by an access command received at step 360. If the Memory Space 205 to be accessed is implemented as a stack, the memory location address computation information includes the memory section information for Memory Section 220 stored to AU 128 at step 420. The memory location address computation information also includes THD# 202 identifying the thread producing the access command, THD# 202 being contained in the address request information received at step 360. If the Memory Space 205 to be accessed is implemented to support random access, the memory location address computation information further includes a memory location offset, the memory location offset being contained in the address request information received at step 360.

[0087] At step 470, AU 128 applies a predefined computation to the memory location address computation information to determine a memory location address of a Memory Location 210 in a Memory Space 205 to be accessed. If the Memory Space 205 to be

accessed is implemented as a stack, the memory location address of Base Memory Space Location 215 is determined using a first predefined computation. The first predefined computation may be a following predefined equation:

base memory section address of Memory Section 220 + (THD# 202 * first memory space size).

[0088] If the Memory Space 205 to be accessed is implemented to support random access, any Memory Location 210 in Memory Space 205 can be accessed. The memory location address of a Memory Location 210 to be accessed is determined using a second predefined computation. The second predefined computation may be a following predefined equation:

base memory section address of Memory Section 220 + (THD# 202 * first memory space size) + memory location offset.

[0089] For example, as shown in FIG. 2, a base memory space address of Memory Space 205 reserved for a thread having THD# 202 of 1 is determined by the sum of the base memory section address for Memory Section 220 (the memory location address of Base Memory Section Location 225) and 1 (THD# 202) multiplied by S1+1 (first memory space size). If Memory Section 220 is implemented to support random access, the memory location address of Example Memory Location 240 can be determined by the sum of the base memory space address of Memory Space 205 reserved for the thread having THD# 202 of 1 (as determined above) and a memory location offset (equal to 1) associated with Example Memory Location 240.

[0090] In an alternative embodiment, the first predefined computation is a following predefined order of concatenation:

{base memory section address for Memory Section 220, THD# 202, memory location offset of 0}.

[0091] In a further embodiment, the second predefined computation is a following predefined order of concatenation:

{base memory section address for Memory Section 220, THD# 202, memory location offset}.

[0092] At step 367, the memory location address of the Memory Location 210 and the operation command is used to determine if a RAW hazard exists. If so, step 367 is repeated until the write operation to the memory location address of the Memory Location 210 is completed. When a RAW hazard does not exist in step 367 the method proceeds to step 370. In a further embodiment, the operation command and address request information is used to determine if a RAW hazard exists. In yet a further embodiment a thread identification number and memory location offset is used to determine if a RAW hazard exists.

[0093] At step 370, the operation command contained in the access command (received at step 360) and the memory location address (determined at step 470) of the Memory Location 210 to be accessed by the access command is sent to Memory Controller 130. If the operation command is a write command, data to be written is also sent to Memory Controller 130. In an alternative embodiment, the stack pointer associated with the thread producing the access command is also sent to Memory Controller 130. Memory Controller 130 then accesses the Memory Location 210 according to the operation command.

[0094] At step 355, the thread continues processing of the sample (received at step 340) according to the program executing on the Functional Unit 140, 150, 160, or 170 that received the sample until the program completes execution. While the thread is

processing the received sample, the Memory Space 205 used by the thread is accessed only by that particular thread.

[0095] The method shown in the flowchart of FIG. 4B can be used to access a Memory Space 205 reserved for another thread in the first set of threads as well. The method of FIG. 4B repeats for each thread in the first set of threads that is assigned (at step 345) to a sample by TCB 127.

[0096] FIG. 5 shows a conceptual diagram of a Memory Resource 200 containing memory spaces that are reserved or managed according to an alternative memory space reservation process. FIG. 5 is described in relation to FIGS. 1 and 2.

[0097] Memory Resource 200 of FIG. 5 contains each element shown in FIG. 2. Memory Resource 200 also contains a Second Memory Section 520 having at least two Memory Spaces 505 for use by threads of a second set of threads executing on Graphics Processor 125. All threads of the first set of threads are of a first thread type and are each reserved a Memory Space 205 having a first memory space size. All threads of the second set of threads are of a second thread type and are each reserved a Memory Space 505 having a second memory space size, the second thread type being different than the first thread type and the second memory space size not being equal to the first memory space size. In the example shown in FIG. 5, the first memory space size is equal to $S1+1$ and the second memory space size is equal to $S2+1$. In an alternate embodiment the second memory space size is equal to the first memory space size.

[0098] Each thread of the first set of threads is identified by an associated first thread type identifier and an associated THD# 202 and each thread of the second set of threads is identified by an associated second thread type identifier and an associated THD# 502. The THD# 202 of a thread in the first set of threads is an order number of the thread in the first set of threads and the THD# 502 of a thread in the second set of threads is an order number of the thread in the second set of threads. In the example shown in FIG. 5, the first set of threads includes $N1+1$ threads (having THD#s 202 ranging from 0 through $N1$) and the second set of threads includes $N2+1$ threads (having THD#s 502 ranging from 0 through $N2$). The number of threads ($N1+1$) in the first set of threads may be equal or not equal to the number of threads ($N2+1$) in the second set of threads.

[0099] A Memory Space 505 is reserved for each thread in the second set of threads. Each Memory Space 505 contains at least one Memory Location 510, each Memory Location 510 being identified by a unique memory location address. A first Memory Location 510 of Memory Space 505 is defined as a Base Memory Space Location 515, the Base Memory Space Location 515 having a unique memory location address (i.e., base memory space address of Memory Space 505).

[00100] Memory Spaces 505 reserved for threads of the second set of threads includes Second Memory Section 520. A first Memory Location 510 of Second Memory Section 520 is defined as a Base Memory Section Location 525, the Base Memory Section Location 525 having a unique memory location address, i.e., a base memory section address of Second Memory Section 520. A size of Second Memory Section 520 is equal to the sum of the memory space sizes of all Memory Spaces 505 allocated to threads of the second set of threads.

[00101] As described above, FIG. 3A shows a flowchart of a method for reserving Memory Spaces 205 for threads of a first set of threads executing on Graphics Processor 125. The method described in relation to FIG. 3A can also be used for reserving Memory Spaces 205 or 505 for threads in a first and second set of threads without substantial alteration. Only those steps that include additional elements or that differ from the steps described above in relation to FIG. 3A will be described in detail.

[00102] At step 305, Driver 113 executes in Host CPU 114 and contains information relating to Graphics Processor 125 including thread type identifiers identifying the type of threads Graphics Processor 125 is capable of executing, a number of threads of each thread type Graphics Processor 125 is capable of executing simultaneously, an amount of memory space used by each thread of each thread type, and a size of memory space available in graphics memory resources located on Graphics Processor 125.

[00103] At step 310, Driver 113 determines a size of Second Memory Section 520 in Memory Resource 200 to be reserved for a second group of threads that will execute on Graphics Processor 125. In the example shown in FIG. 5, the second set of threads includes N_2+1 threads. Driver 113 determines a second memory space size of a Memory Space 505 to be reserved for each thread in the second set of threads, the memory space

size being equal to or greater than the memory space amount used by each thread. In the example shown in FIG. 5, the second memory space size is $S2+1$. Driver 113 then multiplies the number of threads in the second set of threads with the second memory space size to determine the size of Second Memory Section 520 in Memory Resource 200.

[00104] At step 312, Driver 113 determines that Memory Section 220 and Second Memory Section 520 are to be contained in Memory Resource 200. If Memory Resource 200 is a graphics memory resource, at step 315, Driver 113 assigns a base memory section address for Second Memory Section 520. If Memory Resource 200 is Host Memory 112 or a peripheral memory resource (not shown) coupled to System Interface 115, Driver 113 receives, at step 315, a base memory section address for Second Memory Section 520 from Host CPU 114.

[00105] At step 320, Driver 113 allocates a Memory Space 505 to a thread in the second set of threads, starting with a thread having a THD# 502 of 0. If Second Memory Section 520 is implemented as a stack, Driver 113 identifies the base memory space address for Memory Space 505 allocated to a thread by applying the first predefined computation. If Second Memory Section 520 is implemented to support random access, Driver 113 identifies a memory location address for each Memory Location 510 in Memory Space 505 using the second predefined computation. The first and second predefined computations are applied using the base memory section address of Second Memory Section 520, an order number of the thread in the second set of threads (THD# 502), and the second memory space size of a Memory Space 505 allocated to a thread in the second set of threads.

[00106] For example, as shown in FIG. 5, a base memory space address of Memory Space 505 reserved for a thread having THD# 502 of 1 is determined by the sum of the base memory section address for Second Memory Section 520 and 1 (THD# 502) multiplied by $S2+1$ (second memory space size). If Second Memory Section 520 is implemented to support random access, the memory location address of Example Memory Location 540 can be determined by the sum of the base memory space address of Memory Space 505 reserved for the thread having THD# 502 of 1 (as determined

above) and a memory location offset (equal to 2) associated with Example Memory Location 540.

[00107] At step 325, the base memory space address of Memory Space 205 allocated to a thread in the first set of threads, the THD# 202 of the thread, and a thread type identifier associated with the thread are added to a look-up table. Likewise, the base memory space address of Memory Space 505 allocated to a thread in the second set of threads, the THD# 502 of the thread, and a thread type identifier associated with the thread are added to the look-up table. After steps 320 through 330 are completed, a Memory Space 505 is allocated for each thread in the second set of threads and each Memory Space 505 is added to the look-up table. In the example shown in FIG. 5, $N2 + 1$ threads in the second set of threads are allocated a Memory Space 505, each Memory Space 505 being reserved for the thread to which it is allocated.

[00108] As described above, FIG. 3B shows a flowchart of a method for accessing Memory Spaces 205 that have been reserved for threads of a first set of threads using the method shown in FIG. 3A. The method described in relation to FIG. 3B can also be used for accessing Memory Spaces 205 or 505 that have been reserved for threads of a first and second set of threads without substantial alteration. Only those steps that include additional elements or that differ from the steps described above in relation to FIG. 3B will be described in detail.

[00109] At step 340, a Functional Unit 140, 150, 160, or 170 receives a sample, a sample type identifier associated with the sample, and a pointer to a program associated with the sample. At step 345, TCB 127 assigns a thread to the sample based on the thread type identifier associated with the thread and the sample type identifier associated with the sample.

[00110] At step 360, an access command produced by a thread is received at AU 128 including an operation command and address request information. If the Memory Space 505 to be accessed is implemented as a stack, the address request information includes the thread type identifier associated with the thread and THD# 202 or 502 that identifies the thread producing the access command. If the Memory Space 205 or 505 to be accessed is implemented to support random access, the address request information also

includes a memory location offset provided by the instruction generating the access command.

[00111] At step 365, AU 128 uses the look-up table and the received address request information to determine a memory location address of a Memory Location 210 or 510 in the Memory Space 205 or 505 to be accessed. If the Memory Space 205 or 505 to be accessed is implemented as a stack, AU 128 uses the thread type identifier and THD# 202 or 502 contained in the address request information to determine the memory location address of Base Memory Space Location 215 or 515 from the look-up table. If the Memory Space 205 or 505 to be accessed is implemented to support random access, AU 128 uses the thread type identifier, THD# 202 or 502, and the memory location offset contained in the address request information to determine the memory location address of the Memory Location 210 or 510 to be accessed from the look-up table.

[00112] As described above, FIGS. 4A and 4B show flowcharts of a method for managing Memory Spaces 505 for use by threads of a first set of threads executing on Graphics Processor 125. The method described in relation to FIGS. 4A and 4B can also be used for managing Memory Spaces 205 or 505 for use by threads of a first and second set of threads executing on Graphics Processor 125 without substantial alteration. Only those steps that include additional elements or that differ from the steps described above in relation to FIGS. 4A and 4B will be described in detail.

[00113] At step 420, the memory section information for Second Memory Section 520 (i.e., the second memory space size of a Memory Space 505 in Second Memory Section 520 and the base memory section address for Second Memory Section 520) is stored to AU 128.

[00114] At step 465, memory location address computation information is loaded into the AU computational unit of AU 128. If the thread type identifier contained in the address request information (received at step 360) is the first thread type identifier (indicating a thread of the first set of threads), the memory location address computation information includes the memory section information for Memory Section 220. If the thread type identifier contained in the address request information (received at step 360) is the second thread type identifier (indicating a thread of the second set of threads), the

memory location address computation information includes the memory section information for Second Memory Section 520. The memory location address computation information also includes THD# 202 or 502 (contained in the address request information received at step 360). If the Memory Space 205 or 505 to be accessed is implemented to support random access, memory location address computation information further includes a memory location offset (contained in the address request information received at step 360).

[00115] At step 470, AU 128 applies a predefined computation to the memory location address computation information loaded onto the AU computational unit (at step 465) to determine a memory location address of a Memory Location 210 or 510 in a Memory Space 205 or 505 to be accessed. If the Memory Space 205 or 505 to be accessed is implemented as a stack, AU 128 applies the first predefined computation using the memory location address computation information to determine the memory location address of Base Memory Space Location 215 or 515 (i.e., the base memory space address for Memory Space 205 or 505). If the Memory Space 205 or 505 to be accessed is implemented to support random access, AU 128 applies the second predefined computation using the memory location address computation information to determine the memory location address of the Memory Location 210 or 510 to be accessed.

[00116] FIG. 6 shows a conceptual diagram of a Memory Resource 200 containing memory spaces that are reserved or managed according to a further alternative memory space reservation process. FIG. 6 is described in relation to FIGS. 1 and 2.

[00117] Memory Resource 200 of FIG. 6 contains each element shown in FIG. 2. Memory Section 220 contains Memory Spaces 205 or 605 allocated to threads of a first set of threads. The first set of threads includes a first thread group and a second thread group. All threads of the first thread group are of a first thread type and are each reserved a Memory Space 205 having a first memory space size. All threads of the second thread group are of a second thread type and are each reserved a Memory Space 605 having a second memory space size, the second thread type being different than the first thread type and the second memory space size not being equal to the first memory space size. In

the example shown in FIG. 6, the first memory space size is equal to $S1+1$ and the second memory space size is equal to $S2+1$.

[00118] Each thread of the first thread group is identified by an associated first thread type identifier and an associated THD# 202 and each thread of the second thread group is identified by an associated second thread type identifier and an associated THD# 602.

The THD# 202 of a thread in the first thread group is an order number of the thread in the first thread group and the THD# 602 of a thread in the second thread group is an order number of the thread in the second thread group. In the example shown in FIG. 6, the first thread group includes $N1+1$ threads (having THD#s 202 ranging from 0 through $N1$) and the second thread group includes $N2+1$ threads (having THD#s 602 ranging from 0 through $N2$). The number of threads ($N1+1$) in the first thread group may be equal or not equal to the number of threads ($N2+1$) in the second thread group.

[00119] A Memory Space 605 is reserved for each thread in the second thread group. Each Memory Space 605 contains at least one Memory Location 610, each Memory Location 610 being identified by a unique memory location address. A first Memory Location 610 of Memory Space 605 is defined as a Base Memory Space Location 615, the Base Memory Space Location 615 having a unique memory location address (i.e., base memory space address of Memory Space 605).

[00120] Memory Spaces 205 reserved for threads of the first and second thread group includes Memory Section 220. A last Memory Location 610 of Memory Section 220 is defined as an End Memory Section Location 625, the End Memory Section Location 625 having a unique memory location address, i.e., an end memory section address of Memory Section 220. A size of Memory Section 220 is equal to the sum of the memory space sizes of all Memory Spaces 205 and 605 allocated to threads of the first and second thread group.

[00121] As described above, FIG. 3A shows a flowchart of a method for reserving Memory Spaces 205 for threads of a first set of threads executing on Graphics Processor 125. The method described in relation to FIG. 3A can also be used for reserving Memory Spaces 205 or 605 for threads in a first and second thread group, the first and second thread group including the first set of threads, without substantial alteration. Only those

steps that include additional elements or that differ from the steps described above in relation to FIG. 3A will be described in detail.

[00122] At step 305, Driver 113 executes in Host CPU 114 and contains information relating to Graphics Processor 125 including thread type identifiers identifying the type of threads Graphics Processor 125 is capable of executing, a number of threads of each thread type Graphics Processor 125 is capable of executing simultaneously, an amount of memory space used by each thread of each thread type, and a size of memory space available in graphics memory resources located on Graphics Processor 125.

[00123] At step 310, Driver 113 determines a size of a Memory Section 220 in Memory Resource 200 to be reserved for a first and second group of threads (including the first set of threads) that will execute on Graphics Processor 125. To do so, Driver 113 determines a number of threads in the first thread group and a number of threads in the second thread group. In the example shown in FIG. 6, the first thread group includes $N1+1$ threads and the second thread group includes $N2+1$ threads. Driver 113 determines a first memory space size of a Memory Space 205 to be reserved for each thread in the first thread group and a second memory space size of a Memory Space 605 to be reserved for each thread in the second thread group, the memory space size being equal to or greater than the memory space amount used by each thread. In the example shown in FIG. 6, the first memory space size is $S1+1$ and the second memory space size is $S2+1$. Driver 113 then computes the sum of the number of threads in the first thread group multiplied by the first memory space size and the number of threads in the second thread group multiplied by the second memory space size to determine the size of Memory Section 220 in Memory Resource 200.

[00124] At step 312, Driver 113 determines which Memory Resource 200 is to contain Memory Section 220 and assigns or receives a base memory section address for Memory Section 220. Driver 113 then determines an end memory section address of Memory Section 220 using the base memory section address of Memory Section 220 and the size of Memory Section 220 (determined at step 310). As used herein, the memory section information for Memory Section 220 refers to the first memory space size of a Memory Space 205 in Memory Section 220, the second memory space size of a Memory Space

605 in Memory Section 220, the base memory section address of Memory Section 220, and the end memory section address of Memory Section 220.

[00125] At step 320, Driver 113 allocates a Memory Space 605 to a thread in the second thread group, starting with a thread having a THD# 602 of 0. If Memory Section 220 is implemented as a stack, Driver 113 identifies the base memory space address for Memory Space 605 allocated to a thread by applying a third predefined computation. The third predefined computation may be a following predefined equation:

end memory section address of Memory Section 220 - (THD# 602 * second memory space size).

[00126] If Memory Section 220 is implemented to support random access, Driver 113 identifies a memory location address for each Memory Location 610 in Memory Space 605 using a fourth predefined computation. The fourth predefined computation may be a following predefined equation:

end memory section address of Memory Section 220 - (THD# 602 * second memory space size) – memory location offset.

[00127] For example, as shown in FIG. 6, an end memory space address of Memory Space 605 reserved for a thread having THD# 602 of 1 is determined by the end memory section address for Memory Section 220 minus 1 (THD# 602) multiplied by S2+1 (second memory space size). If Memory Section 220 is implemented to support random access, the memory location address of Example Memory Location 640 can be determined by the base memory space address of Memory Space 605 reserved for the thread having THD# 602 of 1 (as determined above) minus a memory location offset (equal to 2) associated with Example Memory Location 640.

[00128] At step 325, the base memory space address of Memory Space 205 allocated to a thread in the first thread group, the THD# 202 of the thread, and a thread type identifier associated with the thread are added to a look-up table. Likewise, the base memory space address of Memory Space 605 allocated to a thread in the second thread group, the THD# 602 of the thread, and a thread type identifier associated with the thread are added to the look-up table. After steps 320 through 330 are completed, a Memory Space 205 or 605 is allocated for each thread in the first and second thread groups and each Memory Space 205 or 605 is added to the look-up table.

[00129] As described above, FIG. 3B shows a flowchart of a method for accessing Memory Spaces 205 that have been reserved for threads of a first set of threads using the method shown in FIG. 3A. The method described in relation to FIG. 3B can also be used for accessing Memory Spaces 205 or 605 that have been reserved for threads of a first and second thread group (including the first set of threads) without substantial alteration. Only those steps that include additional elements or that differ from the steps described above in relation to FIG. 3B will be described in detail.

[00130] At step 340, a Functional Unit 140, 150, 160, or 170 receives a sample, a sample type identifier associated with the sample, and a pointer to a program associated with the sample. At step 345, TCB 127 assigns a thread to the sample based on the thread type identifier associated with the thread and the sample type identifier associated with the sample.

[00131] At step 360, an access command produced by a thread is received at AU 128 including an operation command and address request information. If the Memory Space 205 or 605 to be accessed is implemented as a stack, the address request information includes the thread type identifier associated with the thread and THD# 202 or 602 that identifies the thread producing the access command. If the Memory Space 205 or 605 to be accessed is implemented to support random access, the address request information also includes a memory location offset provided by the instruction generating the access command.

[00132] At step 365, AU 128 uses the look-up table and the received address request information to determine a memory location address of a Memory Location 210 or 610 in

the Memory Space 205 or 605 to be accessed. If the Memory Space 205 or 605 to be accessed is implemented as a stack, AU 128 uses the thread type identifier and THD# 202 or 602 contained in the address request information to determine the memory location address of Base Memory Space Location 215 or 615 from the look-up table. If the Memory Space 205 or 605 to be accessed is implemented to support random access, AU 128 uses the thread type identifier, THD# 202 or 602, and the memory location offset contained in the address request information to determine the memory location address of the Memory Location 210 or 610 to be accessed from the look-up table.

[00133] As described above, FIGS. 4A and 4B show flowcharts of a method for managing Memory Spaces 605 for use by threads of a first set of threads executing on Graphics Processor 125. The method described in relation to FIGS. 4A and 4B can also be used for managing Memory Spaces 205 or 605 for use by threads of a first and second thread group (including the first set of threads) executing on Graphics Processor 125 without substantial alteration. Only those steps that include additional elements or that differ from the steps described above in relation to FIGS. 4A and 4B will be described in detail.

[00134] At step 420, the memory section information for Memory Section 220 (i.e., the first memory space size of Memory Space 205, the second memory space size of Memory Space 605, the base memory section address, and the end memory section address) is stored to AU 128.

[00135] At step 465, memory location address computation information is loaded into the AU computational unit of AU 128. If the thread type identifier contained in the address request information (received at step 360) is the first thread type identifier (indicating a thread of the first thread group), the memory location address computation information includes the first memory space size of Memory Space 205 and the base memory section address of Memory Section 220. If the thread type identifier contained in the address request information (received at step 360) is the second thread type identifier (indicating a thread of the second thread group), the memory location address computation information includes the second memory space size of Memory Space 605 and the end memory section address of Memory Section 220. The memory location

address computation information also includes THD# 202 or 602 (contained in the address request information received at step 360). If the Memory Space 205 or 605 to be accessed is implemented to support random access, memory location address computation information further includes a memory location offset (contained in the address request information received at step 360).

[00136] At step 470, AU 128 applies a predefined computation to the memory location address computation information loaded onto the AU computational unit (at step 465) to determine a memory location address of a Memory Location 210 or 610 in a Memory Space 205 or 605 to be accessed. If the Memory Space 205 or 605 to be accessed is implemented as a stack, the memory location address of Base Memory Space Location 215 or 615 is to be determined. If the thread type identifier contained in the address request information (received at step 360) is the first thread type identifier (indicating a thread of the first thread group), AU 128 applies the first predefined computation using the memory location address computation information to determine the memory location address of Base Memory Space Location 215 (i.e., the base memory space address for Memory Space 205). If the thread type identifier contained in the address request information (received at step 360) is the second thread type identifier (indicating a thread of the second thread group), AU 128 applies the third predefined computation using the memory location address computation information to determine the memory location address of Base Memory Space Location 615 (i.e., the base memory space address for Memory Space 605).

[00137] At step 470, if the Memory Space 205 or 605 to be accessed is implemented to support random access, any Memory Location 210 or 610 in Memory Space 205 or 605 can be accessed. If the thread type identifier contained in the address request information (received at step 360) is the first thread type identifier (indicating a thread of the first thread group), AU 128 applies the second predefined computation using the memory location address computation information to determine the memory location address of Memory Location 210 to be accessed. If the thread type identifier contained in the address request information (received at step 360) is the second thread type identifier (indicating a thread of the second thread group), AU 128 applies the fourth predefined

computation using the memory location address computation information to determine the memory location address of Memory Location 610 to be accessed.

[00138] FIG. 7 shows a conceptual diagram of a Memory Resource 200 containing memory spaces that are reserved or managed according to a yet a further alternative method of the invention. FIG. 7 is described in relation to FIGS. 1 and 2.

[00139] FIG. 7 shows Memory Locations 210 of a Memory Space 205 allocated to a thread of the first set of threads interleaved with Memory Locations 210 of at least one other Memory Space 205 allocated to another thread of the first set of threads to form Interleaved Memory Spaces 705 containing Interleaved Memory Locations 710. In the example shown in FIG. 7, $S1+1$ Memory Locations 210 of $N1+1$ Memory Spaces 205 are interleaved forming $N1+1$ Interleaved Memory Spaces 705 each containing $S1+1$ Interleaved Memory Locations 710. Interleaving the memory spaces for a set of threads can result in lower latency for memory requests for threads accessing neighboring memory locations. In one embodiment, when memory location 0 for THD#0 is read by AU 128 from Local Memory 135 or Host Memory 112, additional memory locations, such as memory location 0 for THD#1 and THD#2 are read and stored in Cache 129. When the threads are being processed simultaneously, memory location 0 for THD#1 and THD#2 can be received from Cache 129 instead of from Local Memory 135 or Host Memory 112.

[00140] As described above, FIG. 3A shows a flowchart of a method for reserving Memory Spaces 205 for threads of a first set of threads executing on Graphics Processor 125. The method described in relation to FIG. 3A can also be used for reserving Interleaved Memory Spaces 705 for threads of a first set of threads without substantial alteration. Only those steps that include additional elements or that differ from the steps described above in relation to FIG. 3A will be described in detail.

[00141] At step 320, Driver 113 allocates an Interleaved Memory Space 705 to a thread in the first set of threads, starting with a thread having a THD# 202 of 0. Driver 113 identifies a memory location address for each Interleaved Memory Location 710 in Interleaved Memory Space 705 using a third predefined computation. The third predefined computation may be a following predefined equation:

base memory section address of Memory Section 220 + (memory location offset * first memory space size) + THD# 202.

[00142] In an alternative embodiment, the third predefined computation is a following predefined order of concatenation that may be used to access data stored in a conventional multi-dimensional array:

{base memory section address for Memory Section 220, memory location offset, THD# 202}.

[00143] At step 325, the THD# 202 of the thread allocated Interleaved Memory Space 705 (at step 320), a memory location offset associated with each Interleaved Memory Location 710 in the allocated Interleaved Memory Space 705, and a memory location address for each Interleaved Memory Location 710 are added to a look-up table.

[00144] As described above, FIG. 3B shows a flowchart of a method for accessing Memory Spaces 205 that have been reserved for threads of a first set of threads using the method shown in FIG. 3A. The method described in relation to FIG. 3B can also be used for accessing Interleaved Memory Spaces 705 that have been reserved for threads of a first set of threads without substantial alteration. Only those steps that include additional elements or that differ from the steps described above in relation to FIG. 3B will be described in detail.

[00145] At step 360, an access command produced by the thread is received at AU 128 containing address request information. The address request information includes the THD# 202 that identifies the thread producing the access command and a memory location offset provided by the instruction generating the access command.

[00146] As described above, FIGS. 4A and 4B show flowcharts of a method for managing Memory Spaces 505 for use by threads of a first set of threads executing on

Graphics Processor 125. The method described in relation to FIGS. 4A and 4B can also be used for managing Interleaved Memory Spaces 705 for use by threads of a first set of threads executing on Graphics Processor 125 without substantial alteration. Only those steps that include additional elements or that differ from the steps described above in relation to FIGS. 4A and 4B will be described in detail.

[00147] At step 465, memory location address computation information is loaded into the AU computational unit of AU 128. The memory location address computation information includes the memory section information for Memory Section 220, THD# 202 identifying the thread producing the access command, and a memory location offset. At step 470, AU 128 applies the third predefined computation to the memory location address computation information to determine a memory location address of an Interleaved Memory Location 710 in an Interleaved Memory Space 705 to be accessed.

[00148] FIG. 8 shows a conceptual diagram of a Memory Resource 200 for storing thread data from a Primary Memory Resource 800. FIG. 8 is described in relation to FIGS. 1 and 2.

[00149] Primary Memory Resource 800 can be Host Memory 112, a peripheral memory resource (not shown) coupled to System Interface 115 (e.g., hard disk drive, Zip drive, tape drive, CD-R, CD-RW, etc.), a graphics memory resource (such as Local Memory 135 or LSR 145, 155, 165, or 175), or any combination of the above. Primary Memory Resource 800 stores thread data generated by at least one thread of the first set of threads executing on Graphics Processor 125. As shown in FIG. 8, Memory Resource 200 may be used to store overflow thread data (spillage) from Primary Memory Resource 800. Primary Memory Resource 800 may or may not be located on the same physical device as Memory Resource 200. Primary Memory Resource 800 may be implemented as a stack or to support random access.

[00150] In one embodiment, as shown in the example of FIG. 8, Primary Memory Resource 800 contains a Primary Memory Space 805 for each thread in the first set of threads that is allocated, accessed, or managed (as described in FIGS. 3A, 3B, 4A, and 4B). As such, each thread in the first set of threads is identified by a THD# 802 and has a reserved Primary Memory Space 805 (having a base memory space address) in a Primary

Memory Section 820 (having a base memory section address). Each Primary Memory Space 805 has at least one Primary Memory Location 810 and a memory space size equal to the number of Primary Memory Locations 810 contained in Primary Memory Space 805.

[00151] In an alternative embodiment, Primary Memory Resource 800 contains a Primary Memory Space 805 to store thread data generated by one thread of the first set of threads. In a further embodiment, there is a Primary Memory Resource 800 for each thread in the first set of threads, each Primary Memory Resource 800 containing a Primary Memory Space 805 for use by the thread and being coupled to Memory Resource 200 and Transfer Unit 860. In the further embodiment, all Primary Memory Resources 800 may be located on the same physical device or each Primary Memory Resource 800 may be located on a different physical device. In yet a further embodiment, Primary Memory Resource 800 contains a memory space containing thread data generated by any thread in the first set of threads using any storage structure or method.

[00152] Thread data transfers between Primary Memory Resource 800 and Memory Resource 200 are controlled by Transfer Unit 860 coupled to Primary Memory Resource 800 and Memory Resource 200. Transfer Unit 860 monitors thread data levels in a Primary Memory Space 805 of Primary Memory Resource 800 and performs thread data transfers between a Primary Memory Space 805 of Primary Memory Resource 800 and a Memory Space 205 of Memory Resource 200. Transfer Unit 860 can be internal to Graphics Processor 125 or external to Graphics Processor 125. If more than one Primary Memory Resource 800 is used, a separate Transfer Unit 860 may be implemented for each Primary Memory Resource 800 or one Transfer Unit 860 may be implemented for all Primary Memory Resources 800.

[00153] When Transfer Unit 860 determines that the thread data level in a Primary Memory Space 805 of Primary Memory Resource 800 reaches a first predetermined threshold, Transfer Unit 860 reads a first predetermined amount of thread data from the Primary Memory Space 805 and stores the first predetermined amount of thread data to a Memory Space 205 of Memory Resource 200. When Transfer Unit 860 determines that

the thread data level in a Primary Memory Space 805 of Primary Memory Resource 800 reaches a second predetermined threshold, Transfer Unit 860 reads a second predetermined amount of thread data from a Memory Space 205 of Memory Resource 200 and stores the second predetermined amount of thread data to a Primary Memory Space 805 of Primary Memory Resource 800. Transfer Unit 860 can be hard-wired to perform the monitoring and transferring functions described above or be configured by a computer program product having a computer readable medium that includes computer program instructions to perform the monitoring and transferring functions.

[00154] FIG. 9 shows a flowchart of a method for transferring data between a Primary Memory Space 805 of Primary Memory Resource 800 (shown in FIG. 8) and a Memory Space 205 of Memory Resource 200. FIG. 9 is described in relation to FIGS. 1 and 8.

[00155] The flowchart shown in FIG. 9 assumes that a first Primary Memory Space 805 of Primary Memory Resource 800 is implemented as a stack, is reserved for a first thread in a first set of threads, and is allocated, accessed, or managed in accordance with (as described in FIGS. 3A, 3B, 4A, and 4B) with modifications as described below. The flowchart shown in FIG. 9 also assumes that a first Memory Space 205 of Memory Resource 200 is implemented as a stack, is reserved for the first thread in the first set of threads, and is allocated or managed (as described in FIGS. 3A and 4A) with modifications as described below.

[00156] At step 925, a Functional Unit 140, 150, 160, or 170 of Graphics Processor 125 receives a sample, a sample type identifier associated with the sample, and a pointer to a program associated with the sample. TCB 127 assigns a first thread in a first set of threads executing on Graphics Processor 125 to the sample, the first thread being identified by a THD# 202 and 802. TCB 127 also assigns a stack pointer to the sample, the stack pointer being associated with the thread assigned to the sample. The first thread uses the received pointer to locate the program associated with the sample and loads the program to the Functional Unit 140, 150, 160, or 170 that received the sample. The first thread processes the sample according to the program, the program executing on the Functional Unit 140, 150, 160, or 170 that received the sample.

[00157] At step 930, during the processing of the sample, the first thread produces thread data to be stored (“pushed”) to a first Primary Memory Space 805 reserved for use by the first thread, first Primary Memory Space 805 being located in a Primary Memory Resource 800. During the processing of the sample, the first thread also reads (“pops”) thread data from first Primary Memory Space 805. Thread data is stored to or read from first Primary Memory Space 805 starting from a base Primary Memory Location 810 of first Primary Memory Space 805 (i.e., thread data is “pushed” or “popped” from a “top” of first Primary Memory Space 805). As such, the first thread accesses the “top” of first Primary Memory Space 805 during processing of the received sample.

[00158] At step 935, Transfer Unit 860 determines if a thread data level in first Primary Memory Space 805 has reached a first predetermined threshold. If so, at step 940, Transfer Unit 860 reads a first predetermined amount of thread data from first Primary Memory Space 805. At step 940, thread data is read from a Primary Memory Location 810 of first Primary Memory Space 805 containing a last thread data stored in first Primary Memory Space 805 (i.e., thread data is “popped” from a “bottom” of first Primary Memory Space 805).

[00159] At step 945, Transfer Unit 860 stores the first predetermined amount of thread data from first Primary Memory Space 805 to a first Memory Space 205 reserved for use by the first thread, first Memory Space 205 being located in a Memory Section 220 of Memory Resource 200. Thread data is stored to first Memory Space 205 starting from a base Memory Location 210 of first Memory Space 205 (i.e., thread data is “pushed” to a “top” of first Memory Space 205). To push thread data to the top of first Memory Space 205, Transfer Unit 860 determines a base memory space address of first Memory Space 205 by using a look-up table (as described in relation to FIG. 3A) or a first predefined computation (as described in relation to FIG. 4B). The method then proceeds to step 950.

[00160] In step 935, if Transfer Unit 860 determines that the thread data level in first Primary Memory Space 805 has not reached a first predetermined threshold, the method proceeds to step 950. At step 950, Transfer Unit 860 determines if the thread data level in first Primary Memory Space 805 has reached a second predetermined threshold. If so, at step 955, Transfer Unit 860 reads a second predetermined amount of thread data from

first Memory Space 205. At step 955, thread data is popped from the top of first Memory Space 205. To pop thread data from the top of first Memory Space 205, Transfer Unit 860 determines a base memory space address of first Memory Space 205 by using a look-up table (as described in relation to FIG. 3A) or a first predefined computation (as described in relation to FIG. 4B).

[00161] At step 960, Transfer Unit 860 stores the second predetermined amount of thread data from first Memory Space 205 to first Primary Memory Space 805. At step 960, thread data is stored to a first available Primary Memory Location 810 of first Primary Memory Space 805 ready to store thread data (i.e., thread data is “pushed” to the “bottom” of first Primary Memory Space 805). The method then proceeds to step 930.

[00162] In step 950, if Transfer Unit 860 determines that the thread data level in first Primary Memory Space 805 has not reached a second predetermined threshold, the method proceeds to step 930. At step 930, the thread continues processing of the sample while storing or reading thread data from first Primary Memory Space 805 according to the program associated with the sample until the program completes execution.

[00163] In the method shown in FIG. 9, note that a thread accesses (at step 930) the top of Primary Memory Space 805 during processing of the received sample and Transfer Unit 860 pushes (at step 960) thread data to the bottom of Primary Memory Space 805 from Memory Space 205. As such, transfer latency from Memory Space 205 to Primary Memory Space 805 can be hidden as long as there is enough thread data in Primary Memory Space 805 for a thread to access without waiting for transfer of thread data from Memory Space 205 to Primary Memory Space 805. The second predetermined threshold (which determines the thread data level of Primary Memory Space 805 when thread data is transferred from Memory Space 205 to Primary Memory Space 805) can be adjusted to achieve minimum transfer latency. Therefore, access time of a thread to Primary Memory Space 805 alone would be the same or substantial similar to the access time of a thread to Primary Memory Space 805 being supplemented by a spillage memory space (Memory Space 205).

[00164] FIG. 10A is a conceptual diagram of a Buffer 1020 displayed by a display device, e.g., monitor, projector, and the like. Data stored in Buffer 1020 is displayed on

Display 185. Additional buffers of arbitrary sizes may be displayed on Display 185. Each buffer may be positioned for display relative to Display 185. A Sample 1040, such as a pixel, within displayed Buffer 1020 is associated with an x,y position relative to Display 185. For example, displayed Buffer 1020 is positioned at an x offset and a y offset relative to the upper left corner of Display 185. The x,y position of Pixel 1040 relative to upper left corner of Display 185 is determined by combining the x offset and y offset with the x,y position of Sample 1040 within displayed Buffer 1020, e.g. relative to the upper left corner of displayed Buffer 1020. The x,y position of Sample 1040 relative to displayed Buffer 1020 is consistent regardless of the position of displayed Buffer 1020 within Display 185. In an alternate embodiment the x,y origin is in the upper left corner of Display 185 and the x,y position of Sample 1040 is described relative to the x,y origin. In this embodiment the x,y position of Sample 1040 changes as the position of displayed Buffer 1020 within Display 185 changes.

[00165] FIG. 10B illustrates a Portion of Graphics Memory 1050, within Memory Resource 200 (shown in FIG. 2) including memory locations storing data for Buffer 1020. Memory locations within a Memory Space 1060 store data for Buffer 1020. For example, a Memory Location 1066 stores data associated with Sample 1040, e.g., color, depth, stencil, shadow depth, map data, and the like. Each sample produced by Graphics Processor 125 uses a predefined amount of memory space for storing fragment data. A size of a Memory Space 1060 (or memory space size) is equal to the number of Memory Locations 1066 contained in Memory Space 1060 multiplied by the size of a Memory Location 1066. The size of a Memory Location 1066 can be any number of bits as specified in hardware or software. A data buffer may include data represented in an 8-bit fixed-point format, a 16-bit fixed-point format, a 16-bit floating-point format, a 32-bit floating-point format, and the like. The number of memory locations contained in a memory space can vary for each data buffer.

[00166] A Memory Location Address 1064 is used to access Memory Location 1066. Memory Location Address 1064 may be computed based on an x,y position within Buffer 1020 and a base memory space address, Memory Location Address 1062, corresponding to a first location within Memory Space 1060. In an alternate embodiment Memory Location Address 1064 is computed based on an x,y position within Display 185, an x

offset of displayed Buffer 1020, a y offset of displayed Buffer 1020, and Location Address 1062. If Memory Space 1060 is implemented as a stack, sample data is accessed from Memory Space 1060 starting from the first Memory Location 1066. In other words, sample data is stored to (“pushed”) and read from (“popped”) a “top” of Memory Space 1060.

[00167] A Memory Space 1070 includes memory locations storing data for another data buffer or memory locations storing data for a thread. Memory Space 1070 has a base memory space address. Each data buffer is associated with a unique data buffer identifier which may be used to determine the data buffer’s corresponding base memory space address.

[00168] Sample data is persistent, the sample data written to a data buffer when a sample is processed by a thread remains in the memory space even after the thread is assigned to another sample. Unlike a memory space allocated to a thread, a memory space allocated to a data buffer is not associated with the thread and is instead accessible by all threads. A memory space used for a data buffer can be located in Host Memory 112, a peripheral memory resource (not shown) coupled to System Interface 115 (e.g., hard drive, Zip drive, tape drive, CD-R, CD-RW, etc.), a graphics memory resource, such as Local Memory 135 or LSR 145, 155, 165, or 175, or any combination of the above. Just as memory spaces in a memory resource for use by threads of Graphics Processor 125 are allocated and managed as previously described, memory spaces for use as data buffers are allocated and managed in accordance with one or more aspects of the invention as described in relation to exemplary embodiments illustratively shown in FIGS. 11A, 11B, and 11C.

[00169] FIG. 11A shows a flowchart for reserving memory spaces for one or more data buffers and threads of a first set of threads executing on Graphics Processor 125. FIG. 11A includes steps described in relation to FIG. 3A. Steps 300 and 305 are performed as described in relation to FIG. 3A. At step 311, Driver 113 determines a size of a memory section in a memory resource to be reserved for a first set of threads that will execute on Graphics Processor 125, as previously described in relation to step 310. Driver 113 also determines a size of one or more memory spaces in the memory resource

to be reserved for data buffers and a unique buffer identifier for each data buffer. Each data buffer is accessible by any thread.

[00170] At step 313, Driver 113 completes the process described in relation to step 312 of FIG. 3A. Driver 113 also determines which Memory Resource 200 is to contain Memory Space 1060. If the size of memory space available in graphics memory resources located on Graphics Processor 125 is equal to or greater than the size of Memory Space 1060 determined at step 311, Driver 113 may determine that a graphics memory resource (such as Local Memory 135 and LSR 145, 155, 165, or 175) is to contain Memory Section 220. In this case, at step 316, Driver 113 assigns a base memory section address for Memory Space 1060. If Driver 113 determines that Host Memory 112 or a peripheral memory resource (not shown) coupled to System Interface 115 is to contain Memory Space 1060, Driver 113 receives, at step 316, a base memory section address for Memory Space 1060 from Host CPU 114. As used herein, the memory space size of Memory Space 1060 and the base memory space address for Memory Space 1060 are referred to as memory space information for Memory Space 1060. The base memory space address for Memory Space 1060 and any other memory spaces allocated to data buffers are added to a data buffer look-up table. In step 316, Driver 113 also assigns a base memory section address for each memory section as described in relation to step 315 of FIG. 3A, allocating a memory space to each data buffer.

[00171] Steps 320, 325, and 330 proceed as described in relation to FIG. 3A. At step 336, the look-up table and the data buffer look-up table are stored to AU 128. In one embodiment the look-up table is combined with the data buffer look-up table.

[00172] In a further embodiment, the information relating to Graphics Processor 125 is stored to AU 128 at step 305. As described above, information relating to Graphics Processor 125 includes a number of threads Graphics Processor 125 is capable of executing simultaneously and a number of data buffers Graphics Processor 125 may access, an amount of memory space used by each thread, and a size of memory space available in graphics memory resources located on Graphics Processor 125. Steps 311

through 330 are then performed by AU 128 instead of Driver 113, as previously described in relation to FIG. 3A.

[00173] FIG. 11B shows a flowchart for accessing memory spaces that have been reserved for one or more data buffers and threads of a first set of threads using the flow shown in FIG. 11A. FIG. 11B includes steps described in relation to FIG. 3B. Steps 340, 345, and 350 are performed as described in relation to FIG. 3B.

[00174] At step 356, the thread processes the sample according to the program associated with the sample, the program executing on the Functional Unit 140, 150, 160, or 170 that received the sample. During the processing of the sample, the thread executes an instruction in the program that generates an access command, such as a write command, for a memory location in Memory Space 1060 allocated to a data buffer. If Memory Space 1060 allocated to the data buffer is implemented as a stack, the access command can also be referred to as a “push” or “pop” command. The allocated Memory Space 1060 to be accessed may be accessed by any thread.

[00175] At step 361, an access command produced by the thread, such as a write command, is received at AU 128. In an alternative embodiment, the stack pointer associated with the thread producing the access command is also received at AU 128. An access command includes an operation command and address request information. Address request information is used by AU 128 in determining a memory location address of a memory location in the memory space to be accessed, such as Memory Location 1066 in the Memory Space 1060. When the memory space to be accessed is allocated to a data buffer, the address request information includes a buffer identifier and position coordinates associated with the sample, e.g., x,y relative to Display 185 or displayed Buffer 1020, provided by the instruction generating the access command. Alternatively, x,y may specify a location within a non-displayable buffer.

[00176] At step 366, AU 128 uses the address request information received in step 361 to determine a memory location address of the memory location in the memory space to be accessed. AU 128 uses the buffer identifier to obtain a base memory space address stored for the data buffer from the look-up table stored to AU 128 (at step 335). AU 128 also uses the position coordinates received by AU 128 (at step 361) to determine the

memory location address of the memory location in the memory space to be accessed by the access command.

[00177] At step 367, the memory location address and the operation command is used by AU 128 to determine if a RAW hazard exists, and, if so, step 367 is repeated until the write operation to the memory location address of the Memory Location 210 is completed. When a RAW hazard does not exist in step 367 the method proceeds to step 370. In a further embodiment, the operation command and address request information received by AU 128 is used to determine if a RAW hazard exists.

[00178] At step 370, the operation command contained in the access command and the memory location address of the memory location to be accessed by the access command are sent to Memory Controller 130. If the operation command is a write command, data to be written is also sent to Memory Controller 130. In an alternative embodiment, the stack pointer associated with the thread producing the access command is also sent to Memory Controller 130. Memory Controller 130 then accesses the Memory Location 210 in Memory Space 1060 according to the operation command.

[00179] At step 356, another thread processes a sample according to the program executing on the Functional Unit 140, 150, 160, or 170 that received the sample. During the processing of the sample, the other thread executes an instruction in the program that generates an access command, such as a read command, for a memory location in Memory Space 1060 allocated to the data buffer.

[00180] At step 361, an access command produced by the other thread, such as a read command, is received at AU 128. Steps 366, 367, and 370 are repeated as previously described with the other thread reading the data buffer which was written by the thread. Furthermore, the memory location in the data buffer written to by the thread may be the same memory location in the data buffer which is read from by the other thread. Therefore, the other thread reads the sample data written by the thread.

[00181] FIG. 11C shows another flowchart for accessing memory spaces that have been reserved for one or more data buffers and threads of a first set of threads using the flow shown in FIG. 11A. FIG. 11C includes steps described in relation to FIG. 11B and FIG. 3B. Steps 340, 345, and 350 are performed as described in relation to FIG. 3B. At

step 352 the thread executes an instruction in the program and determines position coordinates associated with a sample in a data buffer to be written. The position coordinates correspond to a destination location (memory location) in the data buffer. For example, the thread may be executing a scale function which specifies reading sample data from several memory locations, filtering the sample data to produce filtered data, and writing the filtered data to a memory location computed based on a scale factor. Steps 356, 361, 366, 367, and 370 are completed as described in relation to FIG. 11B.

[00182] The invention has been described above with reference to specific embodiments. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention as set forth in the appended claims. The foregoing description and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense. The listing of steps in method claims do not imply performing the steps in any particular order, unless explicitly stated in the claim.